

# Divide-and-Conquer 3D Convex Hulls on the GPU

Jeffrey M. White <sup>\*</sup>      Kevin A. Wortman <sup>†</sup>

March 3, 2013

## Abstract

We describe a pure divide-and-conquer parallel algorithm for computing 3D convex hulls. We implement that algorithm on GPU hardware, and find a significant speedup over comparable CPU implementations.

## 1 Introduction

The *3D convex hull problem* is to identify, for a given set of  $n$  points in  $\mathbb{R}^3$ , the minimal set of input points such that the convex envelope of those points contains all input points. The problem is fundamental to computational geometry and has been studied extensively. Several  $O(n \log n)$  time algorithms are known, with various trade-offs in constant factors, simplicity, numerical robustness, data structure dependencies, and nondegeneracy requirements (see e.g. [2] [4] [7] [8] [10] [13] [15]). Chan’s celebrated output-sensitive algorithm [5] runs in  $O(n \log h)$  time, where  $h$  denotes the number of faces in the output hull, which is asymptotically optimal.

A *graphics processing unit (GPU)* is a parallel coprocessor available in commodity computers. An outgrowth of the computer gaming industry, GPUs utilize a highly-parallel single instruction multiple data (SIMD) architecture. At a high-level, GPUs work by applying a concise constant-space function called a *kernel* to all elements of an array simultaneously. Kernels are written in *domain specific embedded languages (DSEs)* such as NVIDIA’s CUDA [12] or the OpenCL [11] open standard. Each kernel instance is passed an integer *global identifier (id)* which is customarily used to delineate the ranges of input that each kernel invocation applies to. The potential performance, measured in either gigaFLOPS or memory bandwidth, of GPUs is substantially greater than that of multicore CPUs. However, realizing this potential on practical problems, besides the embarrassingly-parallel graphics applications for which GPUs

---

<sup>\*</sup>Department of Computer Science, California State University, Fullerton, [jeffreymarkwhite@gmail.com](mailto:jeffreymarkwhite@gmail.com)

<sup>†</sup>Department of Computer Science, California State University, Fullerton, [kwortman@fullerton.edu](mailto:kwortman@fullerton.edu)

were originally designed, has proven challenging. By and large, existing parallel algorithms depend on facilities, such as message passing and/or synchronization primitives, which are unavailable in the GPU environment. Yet, GPUs are purpose-built for high performance computation on low-dimensional geometric objects, and the opportunity to apply them to computational geometry problems cannot be ignored.

While the 3D convex hull problem has been studied extensively in the standard computational model, precious little past work is applicable to GPU implementations. As stated above, GPU kernels cannot communicate with or synchronize against each other. This limitation rendered unusable every PRAM-model algorithm we surveyed (e.g. [3]). Further, running kernels have no provision for dynamic memory; their collective input and output must be allocated before the kernels execute *en masse* and freed afterward. Accordingly dynamic data structures are off limits. The absence of the doubly connected edge list (DCEL) structure is a particularly formidable obstacle in this context.

There are several results on computing 2D hulls on the GPU [9] [14] [16], but results on the more general and complex 3D problem have been elusive. While preparing this manuscript, we became aware of an independent result on the 3D problem [17]. That algorithm uses heuristics to cull many, but not all, interior points on the GPU, then feeds the remaining points to a black-box CPU hull implementation (e.g. QuickHull [4]). The algorithm presented here achieves competitive performance using a pure GPU divide-and-conquer approach, whose worst case running time is not impacted by the presence of outlier points, and which is conceptually simpler.

## 2 Algorithm

Our algorithm is an adaptation of Chan’s *minimalist* 3D convex hull algorithm [6]. Note that this  $O(n \log n)$ -time algorithm is distinct from the  $O(n \log h)$ -time algorithm mentioned earlier, also authored by Chan. The minimalist algorithm is, by design, a straightforward top-down divide-and-conquer algorithm for computing 3D convex hulls. It was originally motivated by pedagogical needs for an algorithm that achieves a favorable  $O(n \log n)$  running time, while being simple to explain and implement and avoiding dependency on difficult data structures or algorithms. Serendipitously these design constraints correspond to those imposed by the GPU.

The minimalist algorithm works by recasting the 3D problem as a 2D *kinetic* problem. 3D  $(x, y, z)$  points are mapped to  $(x, y, \Delta y)$  points with an initial  $(x, y)$  starting point and  $\Delta y$  vertical rate of speed. As time  $t$  advances, the points move at distinct velocities, which triggers structural changes in the convex hull of the points (see Figure 1). Computing the convex hull of the original 3D points may be visualized as computing a *kinetic movie* of these configurations for all values  $-\infty < t < \infty$ . The algorithm represents this movie as a chronological sequence of *events* when input points are added to, or removed from, the hull. Input points are presorted by  $x$ -coordinate; then event sequences for roughly equal-




Figure 1: Algorithm Events.

size subsets are recursively generated, then combined by a Graham-scan-like  $O(n)$  merging process. In the base case a single point nominates itself as the only convex hull point.

While the minimalist algorithm boasts many of the features necessary for GPU implementation, it cannot be ported to the GPU directly. GPU kernels cannot be recursive, so the top-down divide-and-conquer approach is inappropriate. Instead, the algorithm must be reoriented into one or more mapping steps where an array of input data elements are mapped by a kernel to an array of output data elements. We achieve this reorientation by rewriting the minimalist algorithm to use *bottom-up* divide and conquer. We define a *movie array* data structure as a table of event logs. Our algorithm allocates a single movie array, and initializes one trivial event log for each input point. Then, our algorithm repeats a *merge step* that combines each pair of event logs with adjacent indices into a single event log. A merge step maps a movie array with  $n$  logs of length at most  $l$  to a new array with at most  $\lceil n/2 \rceil$  logs of length at most  $2l$  each. Thus, after  $\lceil \log_2 n \rceil$  merge steps, the movie array contains a single event log for the entire point set. The key property of this algorithm with respect to GPU computation is that each log merge may be performed entirely independently of the others. Each kernel has a particular range of input movie array indices to read from, and a corresponding range of output indices to write to, and may perform its computation independently of other concurrent kernel instances.

### 3 Implementation

Our implementation of the GPU algorithm follows the bottom-up divide-and-conquer design as mentioned above. As shown in Figure 3, the point structure in the CPU algorithm uses a doubly linked list connected by pointers. The idea is to divide the sorted list down into trivial subsequences and build the list back up to the desired set of faces on the convex hull. Memory pointers are difficult (though not impossible) to move between the CPU and GPU since the two devices have distinct memory spaces. Also, on the GPU each kernel instance needs to seek to its assigned sub-input based on its global id, which could take  $O(n)$  time using a list structure. For these reasons, our GPU implementation uses arrayed lists with integer indices rather than linked lists with node addresses (Figure 3).

Modifying the way data is stored impacts the way data is accessed. Figure 4 shows the differences in `act()` function used for inserting and deleting points from event logs. Figure 5 shows the differences in passing potential faces into the event-time calculations.

The implementation process began with converting the original CPU algorithm to use arrays rather than pointers to represent the data. Point data is implemented as its own data type with the  $x$ ,  $y$ , and  $z$  values along with indices to represent the next and previous pointers to reference other points based on their array index. Also, instead of having two pointer lists,  $A$  and  $B$ , we have

```

// CPU Algorithm Point
struct Point {
    double x, y, z;
    Point *prev, *next;
    void act() {...}
};

// GPU Algorithm Point
struct Point {
    cl_float x;
    cl_float y;
    cl_float z;
    cl_int prev;
    cl_int next;
};

```

Figure 2: Differences in the Point datatype.

```

// CPU Algorithm list of points
Point *P = new Point[n];
...
// Sorts points into a doubly
// linked list based x-coordinate.
Point *list = sort(P, n);

// event lists
Point **A = new Point *[2*n];
Point **B = new Point *[2*n];

// GPU Algorithm list of points
Point *P = (Point *)
    malloc(n*sizeof(Point));

// event lists
cl_int *A = (cl_int *)
    malloc(2*n*sizeof(cl_int));
cl_int *B = (cl_int *)
    malloc(2*n*sizeof(cl_int));

```

Figure 3: Differences in list creation.

```

// CPU Algorithm act() function call
point->act()

// CPU Algorithm act() function
struct Point {
...
void act() {
    if (prev->next != this) {
        // insert point
        prev->next = next->prev = this;
    }
    else {
        // delete point
        prev->next = next;
        next->prev = prev;
    }
}
};

// GPU Algorithm act() function call
act(pointIndex);

// GPU Algorithm act() function
void act(int pointIndex) {
    if (P[P[pointIndex].prev].next
        != pointIndex) {
        // insert point
        P[P[pointIndex].prev].next
        = P[P[pointIndex].next].prev
        = pointIndex;
    }
    else {
        // delete point
        P[P[pointIndex].prev].next
        = P[pointIndex].next;
        P[P[pointIndex].next].prev
        = P[pointIndex].prev;
    }
}
}

```

Figure 4: Differences in act() functions.

```

// CPU Algorithm time[0] calculation
t[0] = time(B[i]->prev,
            B[i],
            B[i]->next);

// GPU Algorithm time[0] calculation
t[0] = time(P[B[i]].prev,
            B[i],
            P[B[i]].next);

```

Figure 5: Differences in time calculations.

```

dataOffsetValue = 2;
totalMergesLeft = numberOfPoints/2;
do {
    numberOfThreads = totalMergesLeft;
    runGPUkernels();
    swap(A, B);
    dataOffsetValue = dataOffsetValue*2;
    totalMergesLeft = totalMergesLeft/2;
} while(totalMergesLeft > 1);

```

Figure 6: Main outer loop ran on the CPU to handle the execution of threads on the GPU.

two arrays of indices that reference a master list  $P$  of points.

Another significant change we made to the design is the conversion from a top-down design to a bottom-up design. Instead of using recursion, the heart of the algorithm is placed within one **while** loop as shown in Figure 6. Before implementing this routine as OpenCL kernel code, we wrote a simulation to run on the serial CPU to ensure validity of the algorithm. The ultimate goal of writing a simulation is to avoid the troublesome task of debugging GPU kernel code. This simplified the task of converting the simulation code to GPU kernel code and required only minimal modifications.

Figure 6 shows pseudocode for the main outer loop which runs on the CPU. The main loop uses two movie array structures, both of which exist on the GPU. The two structures alternate between serving as the input and output of a merge step. This approach makes it possible to avoid transferring point data between the GPU and CPU inside the loop, which is desirable as that is an expensive operation. The `dataOffsetValue` is used to calculate the location of where the head of the `leftGroupIndex` and `rightGroupIndex` exist on the globally accessed master list of points  $P$  as shown in Figure 7. To handle the way the CPU algorithm swaps lists  $A$  and  $B$  in each divide routine, we swap the kernel arguments of  $A$  and  $B$  in the `swap(A, B)` function after each iteration of merges. Following the `swap(A, B)` function, `dataOffsetValue` is updated to tie into the next set of group index calculations. Finally, `totalMergesLeft` is cut in half to represent the number threads to take place in the next iteration of merges. When `totalMergesLeft` reaches less than 2, the algorithm exits the main **while** loop as there is no pair of hulls left to be merged together; only one hull is left which represents the final solution.

## 4 Experimental Results

The GPU algorithm shows significant improvements over the CPU algorithm. Peak performance of the GPU algorithm reaches close to a 6x speedup over the CPU algorithm. Figures 8 and 9 illustrate the runtime of both algorithms in milliseconds.

```

// the index of where the head of the
// left group of the list can be found
// on the globally accessed array
leftGroupIndex
= global_ID*dataOffsetValue;

// the index of where the head of the
// right group of the list can be found
// on the globally accessed array
rightGroupIndex
= [leftGroupIndex+((global_ID+1)
*dataOffsetValue)]/2;

// the index of where the globally
// accessed event list begins for the
// group of merges based on the global_ID
eventListOffset = leftGroupIndex*2;

```

Figure 7: GPU kernel code: how the GPU knows which hulls should be merged and which parts of the global data to access.

The CPU algorithm runtime calculations are based on a Intel<sup>®</sup> Core<sup>™</sup> i3-2330M Processor with 2 cores capable of processing 2 threads each at a rate 2.2 gigahertz, and achieves about 25.0 gigaFLOPS according to the LINPACK benchmark tool. The GPU algorithm runtime calculations are based on an ATI Radeon HD 6470m graphics card with 32 stream cores each with 5 processing elements capable of pipelining data at a rate of 700-750 megahertz. Peak performance for this GPU can potentially reach 224-240 gigaFLOPS [1]. Both the CPU and GPU, as described, rank low on the spectrum of hardware available based on performance.

Originally, a hybrid approach to the GPU algorithm seemed to be a more attractive solution to solving the problem. The hybrid GPU algorithm would perform nearly all of the merge steps on the GPU, then perform the last few steps on the CPU after the `totalMergesLeft` variable reached a certain value. The premise of this approach is that the last few iterations are poorly parallelizable and could be more quickly performed by a serial CPU. To accomplish this, the partially computed data would need to be copied from GPU memory to memory that the CPU has access to. On the CPU side, there would be a similar algorithm which would finish the rest of the computation using that same bottom-up style algorithm.

Surprisingly, our experimental results showed that those last few merge iterations take an insignificant amount of time – less than one millisecond. So the hybrid approach is overly-complex, and implementing it would have been an instance of *premature optimization*. The final design of the GPU algorithm takes place entirely on the GPU rather than on both GPU and CPU hardware. The GPU algorithm just requires the use of the CPU for the required OpenCL setup routines and ultimately to read in the data and output the data; the GPU completes all the extensive computations.



Points	GPU Alg. (ms)	CPU Alg. (ms)
4	0	0
8	0	0
16	0	0
32	0	0
64	0	0
128	0	1
256	0	3
512	0	7
1024	0	9
2048	1	10
4096	2	13
8192	4	25
16384	9	34
32768	23	59
65536	38	123
131072	56	233
262144	88	465
524288	173	941
1048576	354	1869
2097152	652	3732
4194304	1309	7494
8388608	2598	15047

Figure 8: Run time for  $n$  data points.




Figure 9: Run time graph for  $n$  data points.

Something we found interesting is the ratio of speedup improvements over the CPU algorithm as the data set increases. For a data set of four points, the speed up is close to 6x. As the data set approaches 32,768 points, the speedup decreases to about 2.5x. From 32,768 points and on, the speedup increases back to about 6x.

Our roughly 6x speedup is notable since it approaches the maximum potential improvement achievable on our hardware. According to the manufacturers, our GPU is capable of roughly 9 times more gigaFLOPS than our CPU. So the greatest conceivable speedup factor is roughly 9, which would correspond to an embarrassingly-parallel problem with negligible overhead. Our implementation comes close to realizing this full potential despite the obstacles inherent in parallelizing the 3D convex hull problem.

## 5 Conclusion

We have shown that bottom-up adaptation of the minimalist divide-and-conquer algorithm for 3D convex hulls is fast, practical, and reasonably straightforward. The approach achieves run-time performance comparable to past GPU implementations of convex hull algorithms.

In performing this exercise, we did make two counterintuitive conclusions. First, while OpenCL and CUDA are intended to be high-level abstractions of GPU hardware, we nonetheless faced many obstacles related to low-level concerns such as memory management, memory hierarchies, and thread scheduling. Second, our intuition was that the overhead of starting and scheduling kernel applications would become a major bottleneck in the later steps of the algorithm. However, empirical results demonstrated this to be a non-issue.

The following are potential areas for future work:

- Higher-level libraries or tools for implementing divide-and-conquer algorithms on the GPU.
- A suite of compatible, parallel GPU implementations of fundamental computational geometry algorithms.
- In particular, an arrangement data structure, e.g. doubly connected edge list, is a prerequisite to implementing many well-motivated algorithms.

## References

- [1] *AMD Accelerated Parallel Processing*. Advanced Micro Devices, Inc., [http://developer.amd.com/sdks/amdappsdk/assets/amd\\_accelerated\\_parallel\\_processing\\_opengl\\_programming\\_guide.pdf](http://developer.amd.com/sdks/amdappsdk/assets/amd_accelerated_parallel_processing_opengl_programming_guide.pdf).
- [2] S. G. Akl and G. T. Toussaint. A fast convex hull algorithm. *Information Processing Letters*, 7(5):219 – 222, 1978.
- [3] N. M. Amato and F. P. Preparata. A time-optimal parallel algorithm for 3D convex hulls. 1993.

- [4] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.*, 22(4):469–483, Dec. 1996.
- [5] T. Chan. Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discrete & Computational Geometry*, 16:361–368, 1996. 10.1007/BF02712873.
- [6] T. M. Chan. A minimalist’s implementation of the 3-d divide-and-conquer convex hull algorithm. 2003.
- [7] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10:377–409, 1993. 10.1007/BF02573985.
- [8] W. F. Eddy. A new convex hull algorithm for planar sets. *ACM Trans. Math. Softw.*, 3(4):398–403, Dec. 1977.
- [9] T. Jurkiewicz and P. Danilewski. Efficient quicksort and 2d convex hull for cuda, and msimd as a realistic model of massively parallel computations. November 2011.
- [10] D. G. Kirkpatrick and R. Seidel. The ultimate planar convex hull algorithm? *SIAM Journal on Computing*, 15(1):287–299, 1986.
- [11] A. Munshi. The OpenCL specification version 1.0.
- [12] NVIDIA. *NVIDIA CUDA Programming Guide 2.0*. 2008.
- [13] F. P. Preparata and S. J. Hong. Convex hulls of finite sets of points in two and three dimensions. *Commun. ACM*, 20(2):87–93, Feb. 1977.
- [14] A. Rueda and L. Ortega. Geometric Algorithms on CUDA. In *Proceedings of the 3<sup>rd</sup> International Conference on Computer Graphics Theory and Applications*, 2008.
- [15] R. Seidel. A convex hull algorithm optimal for point sets in even dimension. Master’s thesis, Dept. of Computer Science, University of British Columbia, Vancouver, Canada, 1981.
- [16] S. Srungarapu, D. Reddy, K. Kothapalli, and P. Narayanan. Fast two dimensional convex hull on the gpu. In *Advanced Information Networking and Applications (WAINA), 2011 IEEE Workshops of International Conference on*, pages 7 –12, march 2011.
- [17] M. Tang, J. yi Zhao, R. Tong, and D. Manocha. GPU accelerated convex hull computation. In *Shape Modeling International (SMI) 2012*, 2012.